



Amendments to the Specification

Paragraph beginning at page 2, line 16

Note also that once the network routing algorithms have converged after a fault, it may be preferable[[.]] to re-optimize the network by performing a reroute based on the current state of the network and network policies in place.

Paragraph beginning at page 2, line 19

In contrast to path rerouting, protection switching recovery mechanisms pre-establish a recovery path or path segment[[.]] based on network routing policies and the restoration requirements of the traffic on the working path. Preferably, the recovery path is link and node disjoint with the working path. When a fault is detected, the protected traffic is switched over to the recovery path(s) and restored.

Paragraph beginning at page 3, line 4

An example of protection switching in MPLS networks is described below. A diagram illustrating an example MPLS based network incorporating a bypass tunnel is shown in Figure 1. The network, generally referenced 10, comprises a plurality of label switched routers (LSRs) 12 connected by links 14. Backup tunnels are established for protecting LSPs statically by the management station or using RSVP signaling. RSVP extensions for [[up]] backup LSP tunnels have been defined. To meet the needs of real-time applications such as video on demand, voice over IP, etc., it is desirable to affect the repair of LSP tunnels within tens of milliseconds. Protection switching can provide such repair times.

Paragraph beginning at page 4, line 16

There is therefore a need for a protection mechanism that is capable of protecting working paths against failures that are not protected by a protection tunnel. Such failures include for example node failures and link failures that [[are]] were not originally protected by a protection tunnel or were originally protected but the protection tunnel is not available.

Paragraph beginning at page 5, line 2

The present invention provides a path reroute mechanism for use in communication networks. The mechanism comprises a multistage search for a routing path to restore traffic following a failure that could not be protected by a previously established protection route (i.e.



protection tunnel, bypass, etc.). The invention is applicable for use in any traffic engineering technique scheme used to create connections in a network. The traffic engineering used must rely on some type of link state information advertising whereby the available bandwidths on each link is flooded throughout the network and used in making bandwidth calculations to insure that connections will provide previously agreed on Committed Information Rates (CIRs).

Paragraph beginning at page 5, line 11

The invention is not limited by the type of network in use, the type of routing used to select the reroutes or by the [[or]] signaling in use in the network. Examples of applicable networks include, but are not limited to, MPLS and Private Network to Network Interface (PNNI) based Asynchronous Transfer Mode (ATM) networks.

Paragraph beginning at page 6, line 8

When needing to route or reroute an LSP, the LSP-initiating node first runs a routing search algorithm such as Dijkstra but only on the links in which the available bandwidth exceeds the committed information rate (CIR) requirement of user LSPs. If this first search ~~did not succeed~~ is not successful, the LSP-initiating node runs Dijkstra a second time but considering the bandwidth reserved for protection purposes as well as that reserved for non-protection purposes. If two types of user LSPs are required (i.e. protected and unprotected LSPs), a three step search is performed.

Paragraph beginning at page 13, line 13

Revertive mode refers to a recovery mode in which traffic is automatically switched back from the recovery path to the original working path upon the restoration of the working path to a fault-free condition. This assumes a failed working path does not automatically surrender resources to the network. Non-revertive mode refers to a recovery mode in which traffic is not automatically switched back to the original working path after this path is restored to a fault-condition. Note that depending on the configuration, the original working path may, upon moving to a fault-free condition, become the recovery path, or it may be used for new working traffic, and [[be]] no longer be associated with its original recovery path.

Paragraph beginning at page 14, line 9

A fault recovery signal (FRS) is a signal that indicates a fault along a working path has been repaired. Again, like the FIS, it is relayed by each intermediate LSR to its upstream or

downstream neighbor, until ~~[[is]]~~ it reaches the LSR that performs recovery of the original path. The FRS is transmitted periodically by the node/nodes closest to the point of failure, for some configurable length of time.

Paragraph beginning at page 14, line 32

For illustration purposes, the principles of the present invention are described in the context of an MPLS based network employing any suitable routing such as OSPF-TE and signaling such as RSVP-TE signaling. Further, it is intended that the mechanism of the present invention is be implemented in ~~[[the]]~~ network devices such as routers (e.g., LSRs in MPLS networks) located within the network. It is not intended, however, that the invention be limited to the configurations and embodiments described herein. It is appreciated that one skilled in the networking, electrical and/or software arts may apply the principles of the present invention to numerous other types of networking devices and network configurations as well without departing from the spirit and scope of the invention.

Paragraph beginning at page 15, line 20

Thus, in accordance with the present invention, the path rerouting mechanism is operative such that the percentage of the bandwidth reserved for protection purposes is only used for rerouting CIR LSPs when no other type of bandwidth is available. This percentage of bandwidth saved for protection purposes is referred to as protection bandwidth.

Paragraph beginning at page 15, line 24

The path reroute mechanism comprises performing multiple searches for a route when a LSP is to be rerouted. Each search is performed on a larger and larger portion of the total available bandwidth of the link. In the example embodiment presented herein, three different bandwidth types are defined including: bandwidth reserved for protection purposes, bandwidth reserved for protected paths (i.e. protected LSPs) and bandwidth reserved for unprotected paths (i.e. unprotected LSPs). Thus, in the first attempt to find a route, only a portion of the available bandwidth of each link is considered. If no route is found, a second search is performed considering a larger portion of the available bandwidth and if a route is still ~~[[nor]]~~ not found, a third search is performed considering 100% of the available bandwidth of the links.

Paragraph beginning at page 16, line 1

Thus, in order for a node to be able to reroute a LSP, it must have knowledge about the availability of the different bandwidths on each link. This is achieved by having each node advertise the actual amount of bandwidth available on each link. Many routing protocols already include such a capability including, for example, OSPF with traffic engineering extensions (OSPF-TE), ATM's PNNI, etc.

Paragraph beginning at page 16, line 23

The mechanism of the present invention is adapted to take into account the existence of multiple traffic priority classes and the overbooking factor used by nodes in allocating bandwidth to LSPs. In many network implementations, the LSRs maintain a different overbooking factor for each priority class. Further, the overbooking factor is local to the LSR and is not necessarily known to the LSP initiating node. Thus, in accordance with the invention, the available bandwidth values advertised by each node already have the overbooking factor factored into them ~~the overbooking factor~~. Thus, the available bandwidth values used by the LSP initiating node in its route search accurately represent the amount of bandwidth available on that link.

Paragraph beginning at page 17, line 6

The structure of the various TLVs is described hereinbelow. Each link in the network is adapted to advertise at least three values: (1) the amount of protection bandwidth available which is the bandwidth available for protection purposes; (2) the amount of bandwidth available for protected LSPs, i.e. LSPs that are protected by a protection tunnel (or tunnels) on other links; and (3) the amount of bandwidth available for unprotected LSPs, i.e. LSPs that are not protected by any protection tunnels. As stated ^{supra} hereinabove, several variations are possible, including using the link-bandwidth sub-TLV provided by the OSPF-TE protocol instead of one of these three TLVs as well as advertising less information.

Paragraph beginning at page 20, line 5

In accordance with the invention, the rerouting algorithm of the present invention is applicable in various cases. One is when a failure or a series of failures is not completely covered by the fast protection mechanism and rerouting is performed in order to fix the LSPs broken due to these failures. A second case is when there is a need to optimize LSPs after a failure, since the fast protection mechanism ~~route~~ routes packets over non-optimized paths. A

third case is in optimizing rerouting after a topology change (e.g., link or node recovery) or as part of a periodic optimization process. When rerouting a working-LSP, such as in the last two cases, the rerouting should not be put into place if the new calculated LSP uses bandwidth allocated for protection tunnels in all or part of its links, while the existing LSP does not. One possible scheme is to prefer the route in which less links are over occupied.